

# **CDMI and Cloud Federation**

**David Slik**  
**NetApp, Inc.**

# Session Agenda

- ❑ A Brief Overview of CDMI
- ❑ What is Federation?
- ❑ Examples of Federations
- ❑ Federation Levels
- ❑ Using CDMI to enable Federation
- ❑ Key CDMI Concepts to enable Federation

## Cloud Data Management Interface



Cloud Storage TWG

**The CDMI standard has been developed over the last year by leading storage vendors, users and researchers of cloud technology**

This session assumes a basic understanding of CDMI concepts and terminology

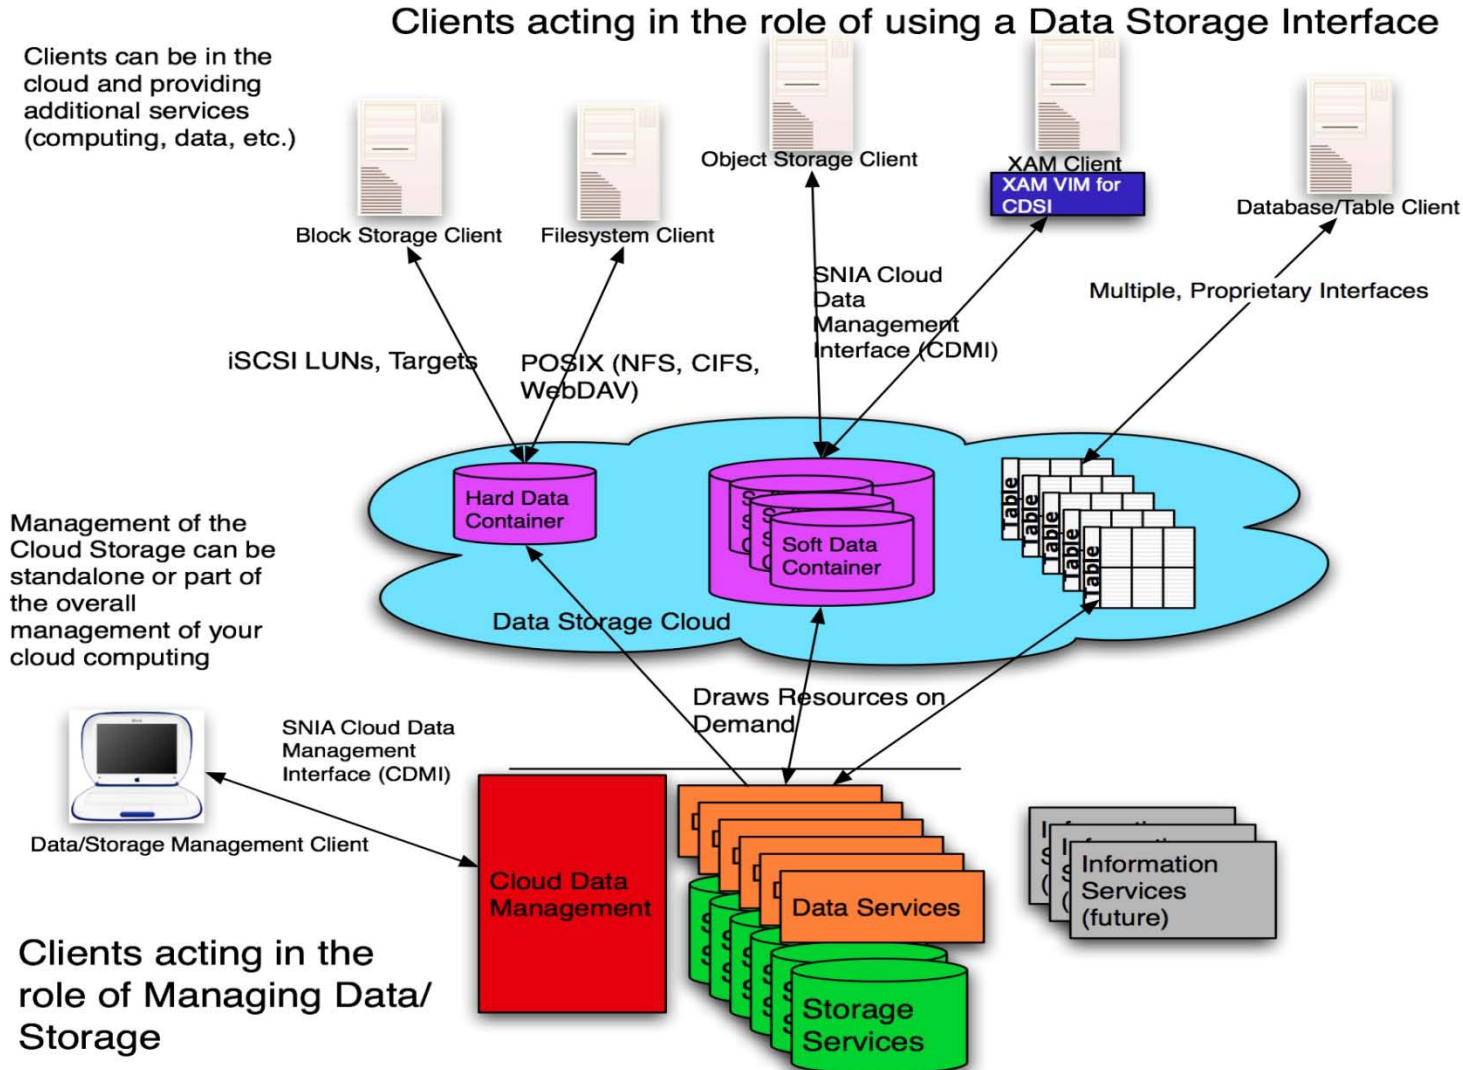
# A Brief Overview of CDMI

- ❑ CDMI has the following goals:
  - ❑ To provide a standard interface for clients to communicate with storage clouds
  - ❑ To provide a standard approach for adding vendor-specific functionality without breaking client compatibility
  - ❑ **To enable standardized Cloud-to-Cloud Federation use cases**

For more details on use cases, see:

[http://www.snia.org/tech\\_activities/publicreview/CloudStorageUseCasesv0.5.pdf](http://www.snia.org/tech_activities/publicreview/CloudStorageUseCasesv0.5.pdf)

# A Brief Overview of CDMI

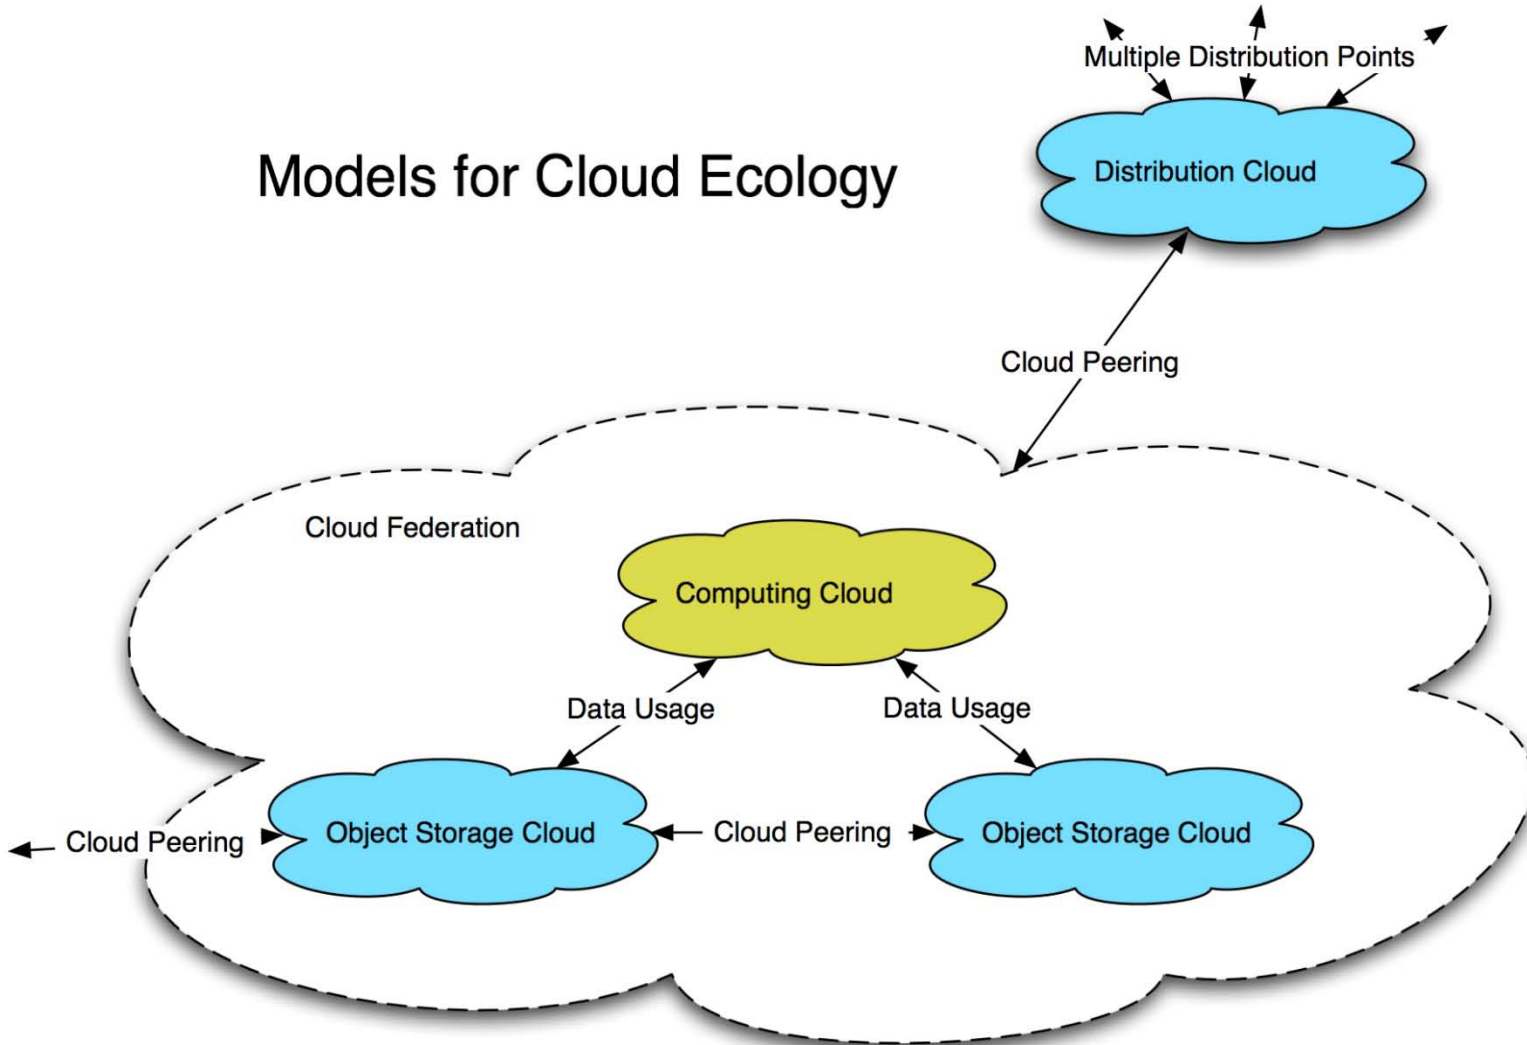


- CDMI provides:
  - A standardized API for client interactions built on top of JSON and RESTful HTTP
  - A standardized object and metadata model for data storage and management
  - A standardized query and notifications model
  - A standardized foundation for multi-tenancy, ownership and **federation**

For more details on the CDMI standard, see:  
<http://www.snia.org/cloud/> and <http://cdmi.sniacloud.com/>

# What is Federation?

## Models for Cloud Ecology



## A Federation is:

A collection of independent entities working together in order to act as a single entity while retaining their individual autonomy.


- There are many uses for federation:
  - Scaling beyond the limits of a single storage system
  - Delegated storage (eg, SSP)
  - Distribution of Stored Data (CDN)
  - Sharing of Stored Data



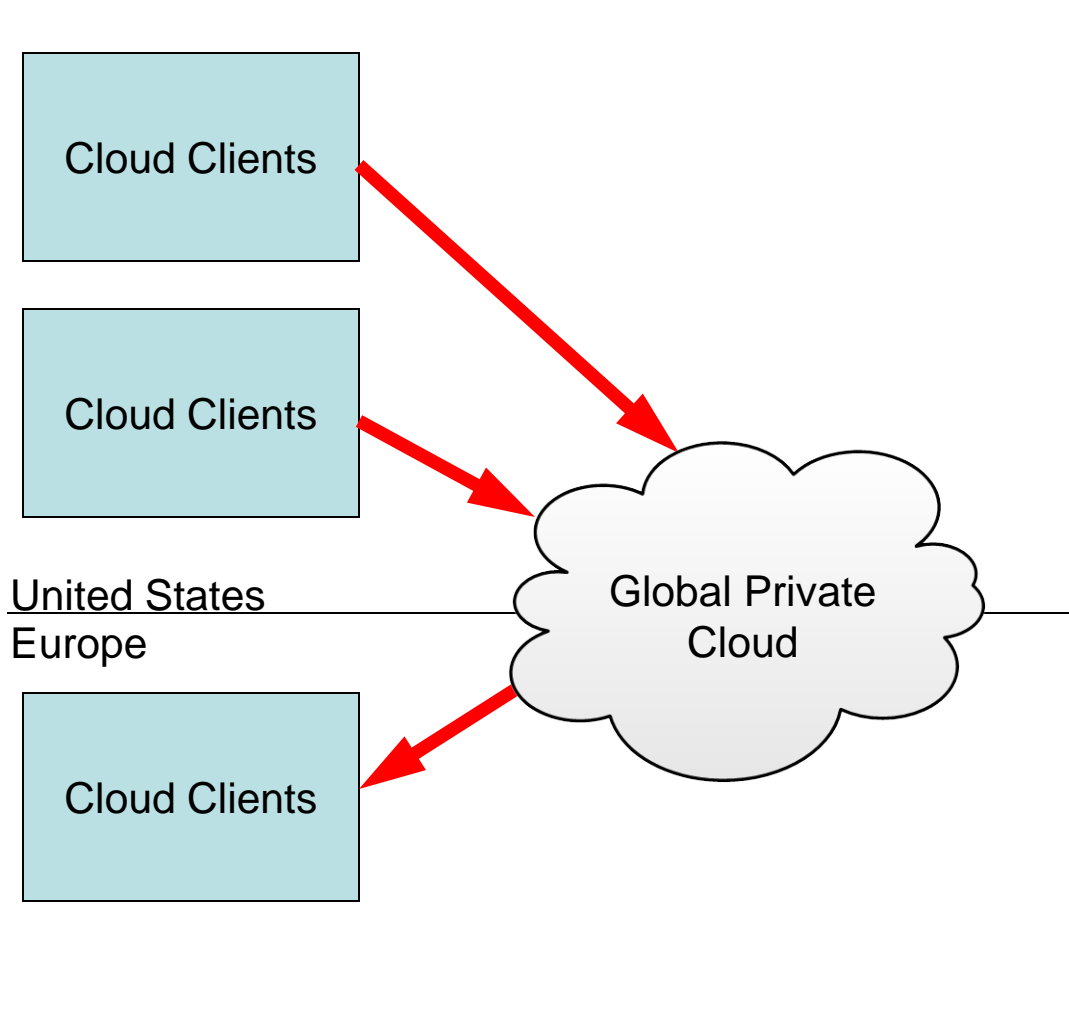
# What is Federation?

- ❑ Storage Federation
  - ❑ Additional storage clouds are used to store objects, but objects are only accessible through the originating cloud
  - ❑ Allows scaling of storage capacity and geographic distribution
- ❑ Access and Peering Federation
  - ❑ Additional storage clouds are used to store objects, and objects can be accessed through any of the federated clouds
  - ❑ Allows scaling of storage capacity, geographic distribution, performance and number of clients
  - ❑ Enables content sharing and distribution (CDN)

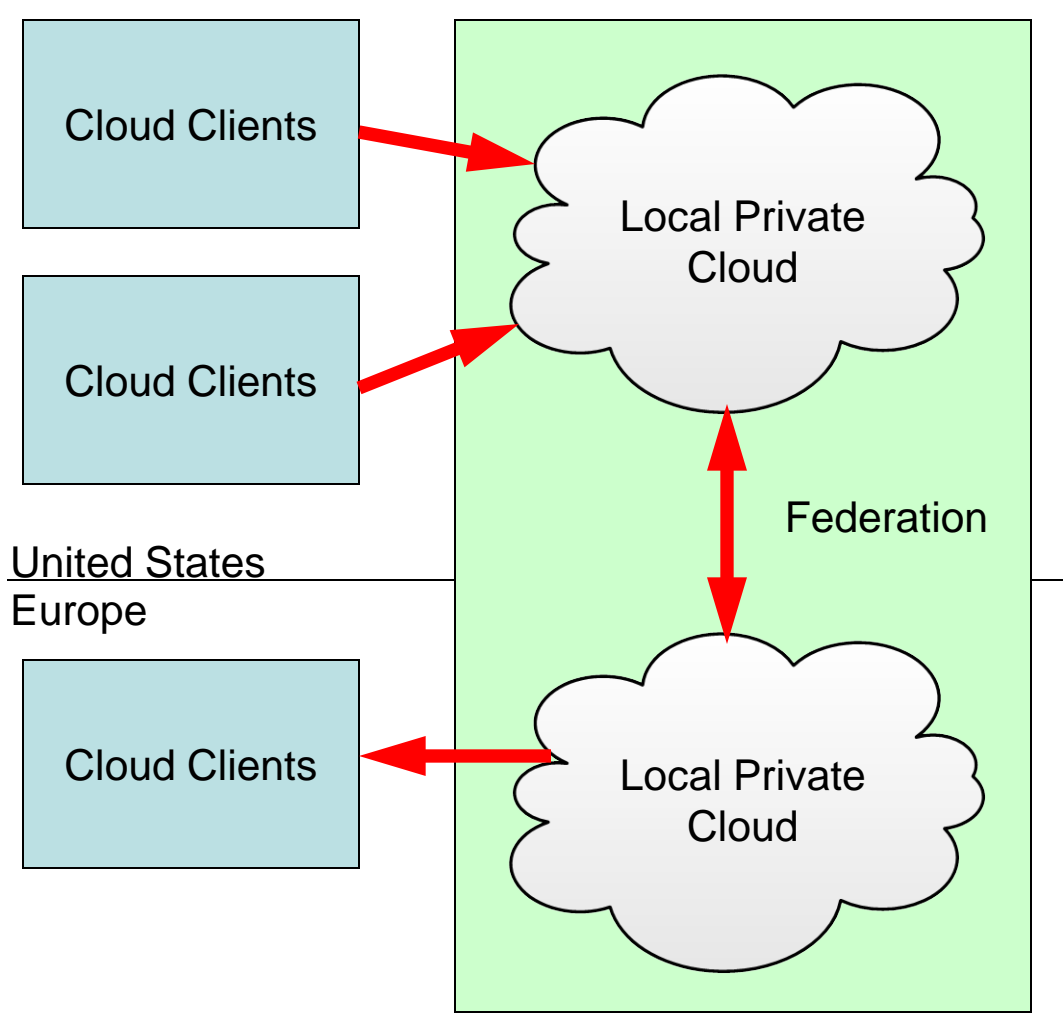
# Non-Federated Private Cloud




# Non-Federated Private Cloud



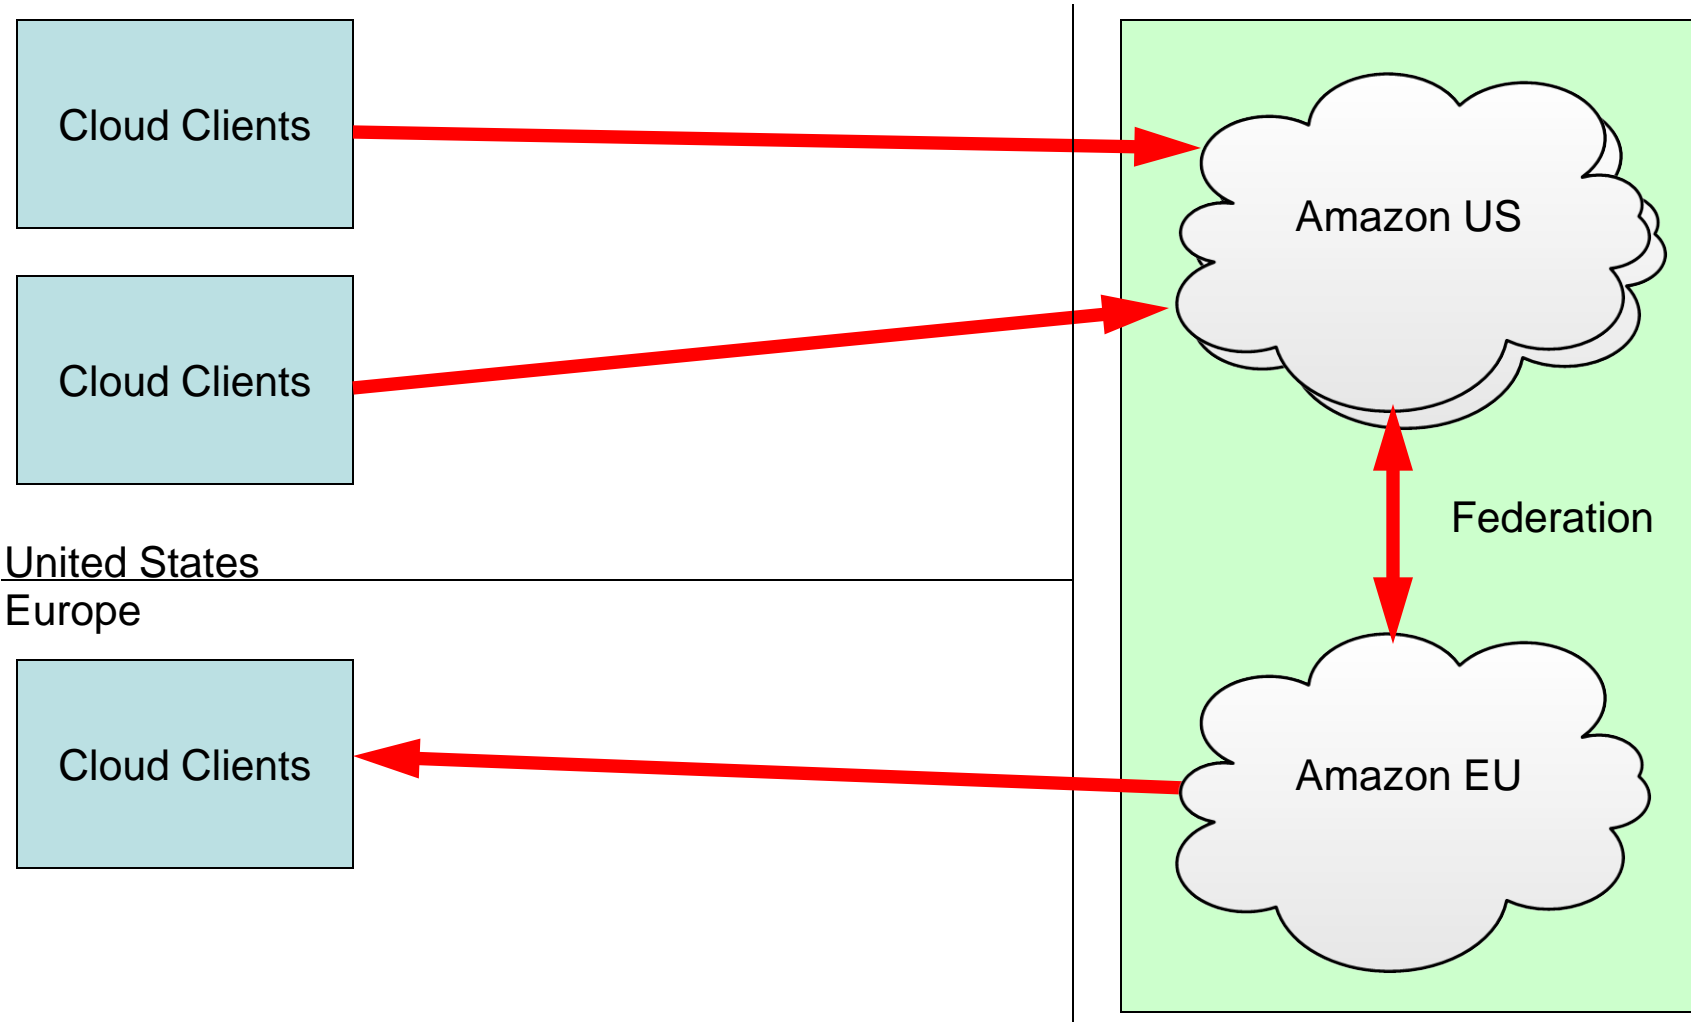
# Federated Private Cloud



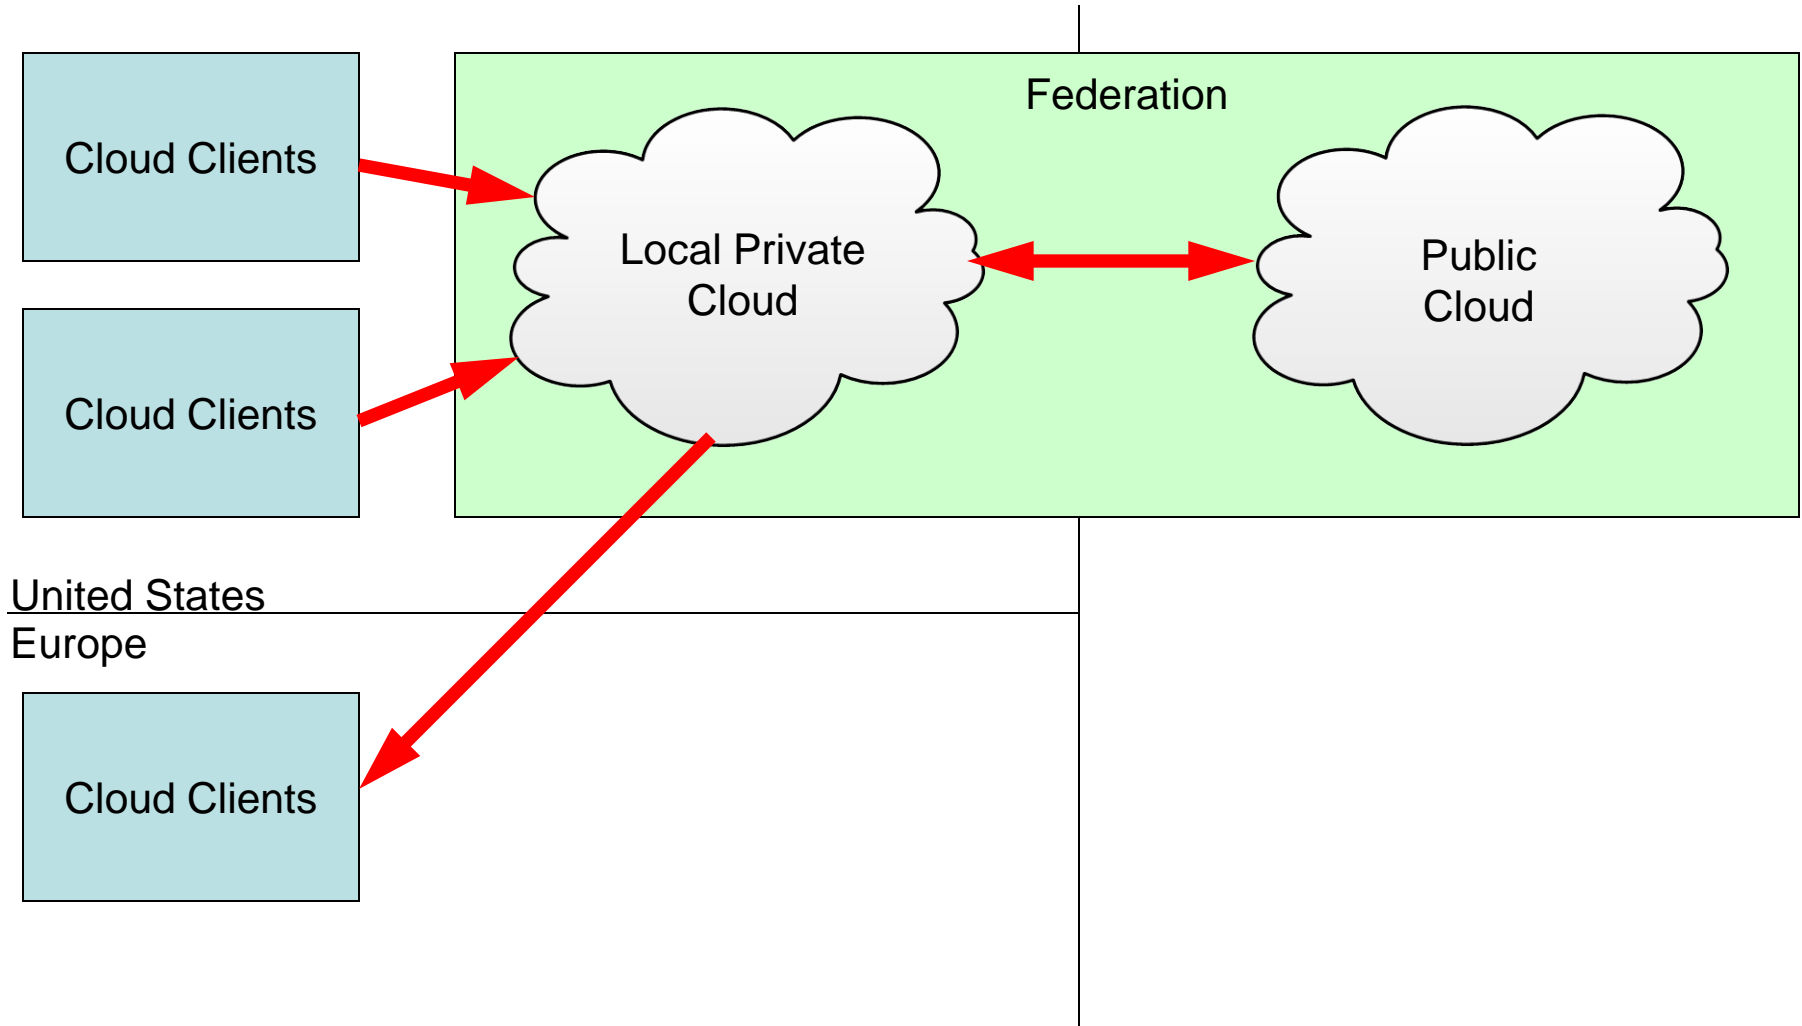
# Non-Federated Public Cloud



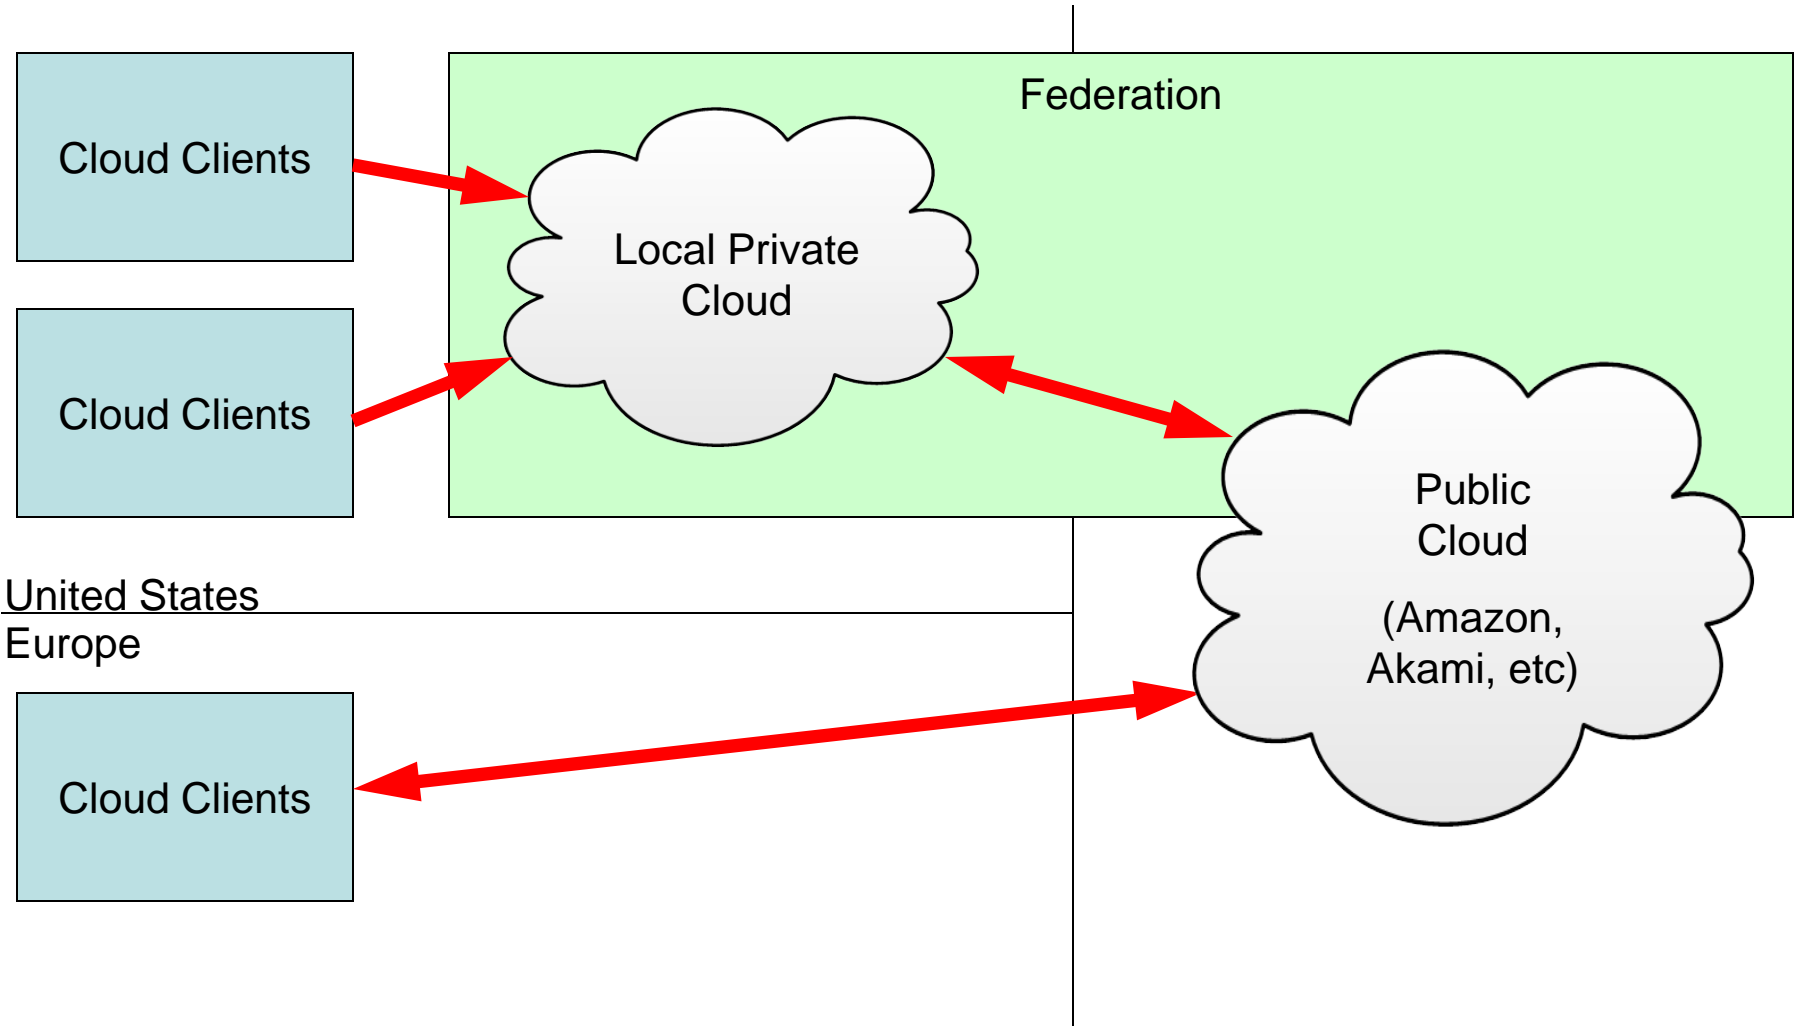
# Federated Public Cloud



# Federated Hybrid Cloud (I - I)

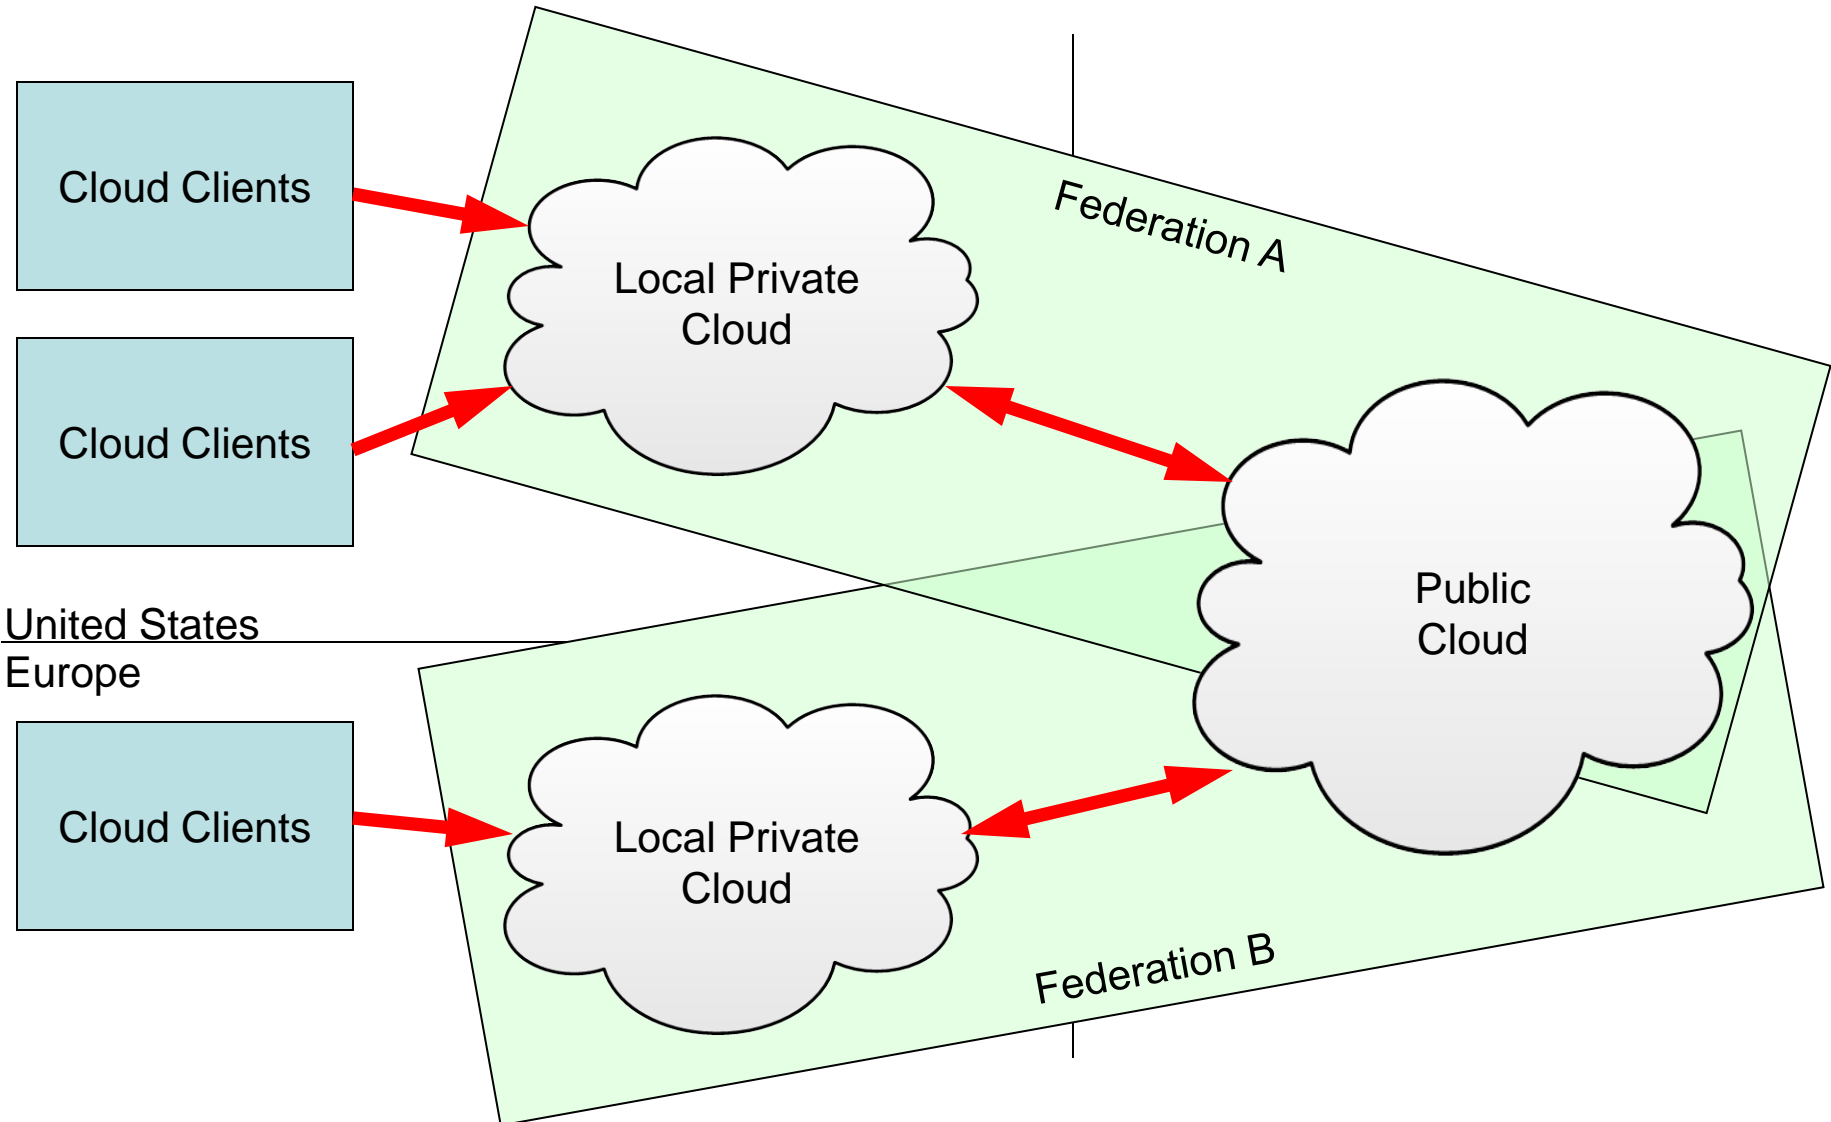


# Federated Hybrid Cloud (I - I)

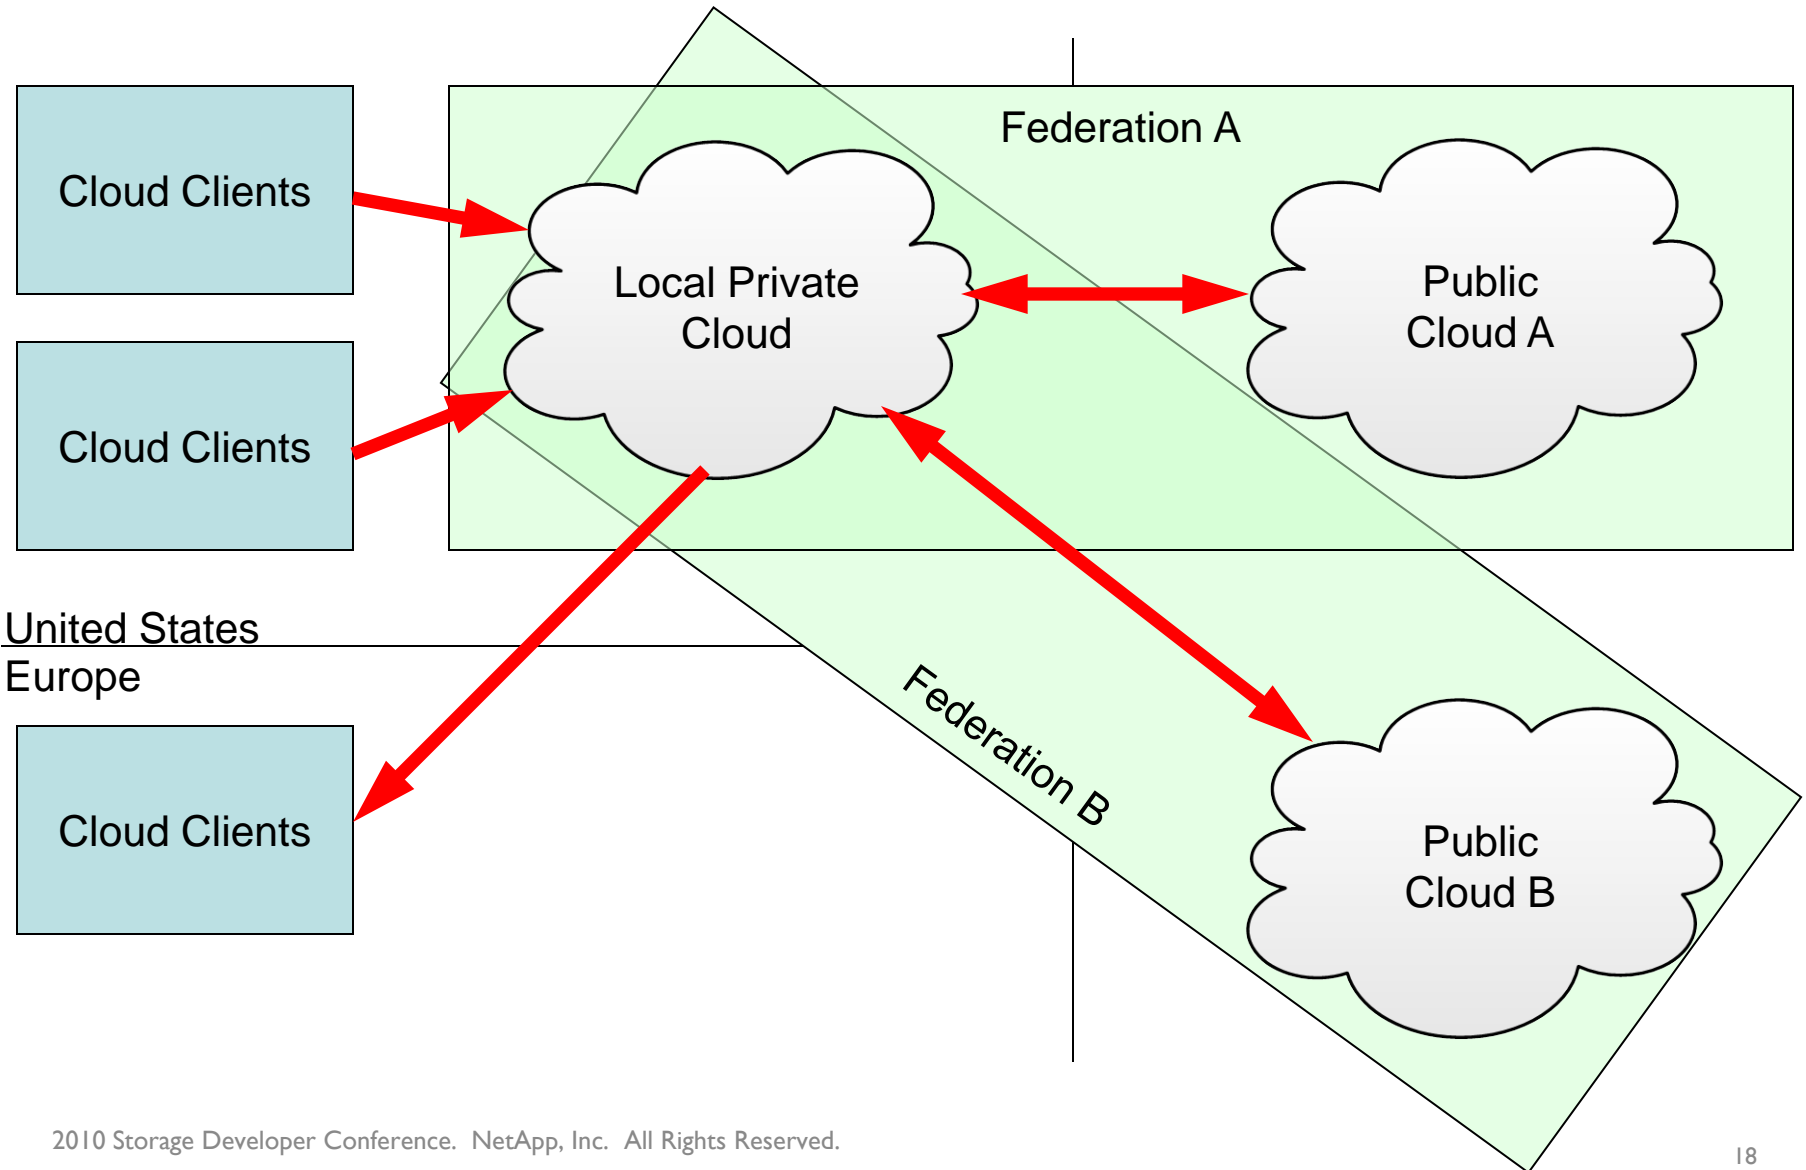




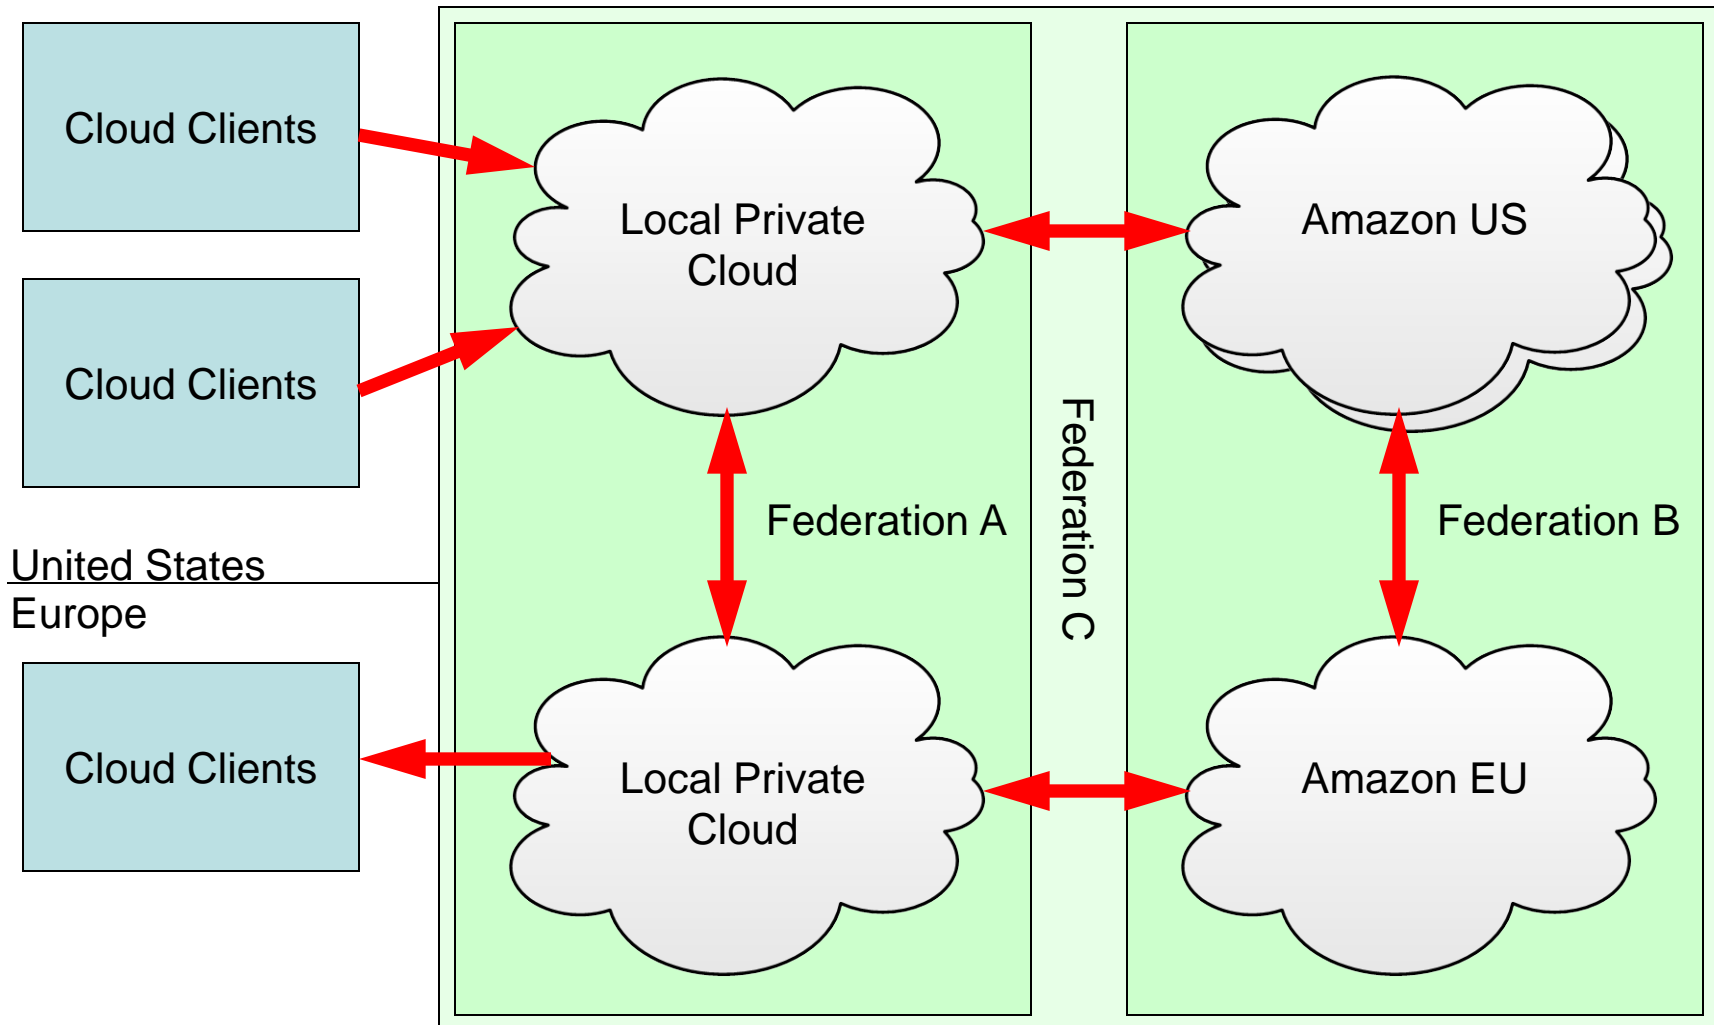
# Federated Hybrid Cloud (2 – 1)



# Federated Hybrid Cloud (1 - 2)



# Federations of Federations (2 – 2)




# Levels of Federation

- ❑ Level 1 – Topology Federation
- ❑ Level 2 – Storage Federation
- ❑ Level 3 – Access Federation
- ❑ Level 4 – Peering Federation

	1	2	3	4
Local cloud mutually authenticated with remote cloud	✓	✓	✓	✓
Local cloud sees remote cloud as a storage destination	✓	✓	✓	✓
Local cloud has private storage area on remote cloud		✓	✓	✓
Local clients can access remote content through local cloud		✓	✓	✓
Local cloud has public storage area on remote cloud			✓	✓
Local clients can access content from remote cloud			✓	✓
Local cloud manages content created on remote cloud				✓
Local clients can access and store content from remote cloud				✓


# Topology Federation (Level I)

- I. Remote cloud administrator defines how the remote cloud looks to the local cloud.
  - Think of this as a “Cloud Export”
  - Like setting up a line of credit with a bank – All the hard work is done up front
  - This is where tenancy management, capacity limits, naming, billing, trust relationships, etc. are defined
  - This process can be automated



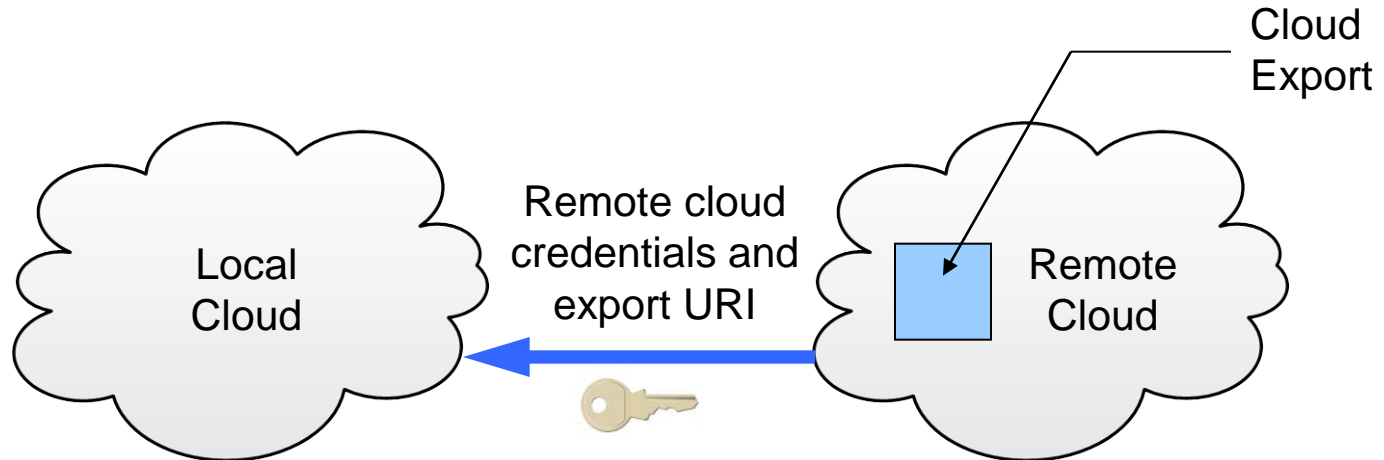
# Topology Federation (Level 1)

2. Credentials and the CDMI domain for the local cloud (as the tenant) are passed to the administrator of the remote cloud. The remote cloud administrator then enrolls the local cloud's credentials to allow access.
  - Once enrolled, the remote cloud has a `cdmi_domain` with the same name as on the local cloud




# Topology Federation (Level 1)

3. Credentials and the cloud export URI are passed back to the administrator of the local cloud. The local cloud administrator then registers the credentials and export with the local cloud.
  - At this point, the local cloud can communicate with the local cloud, and the remote cloud export becomes visible for use in policies.



# Storage Federation (Level 2)

- I. Policies determined by the local cloud administrator indicate that objects should be stored on the remote cloud:
  - ❑ Local cloud does an CDMI PUT (by ID or path) to the URI of the remote cloud export
  - ❑ Local cloud (optionally) records that the object is stored in the remote cloud
  - ❑ Local cloud (optionally) deletes local copy

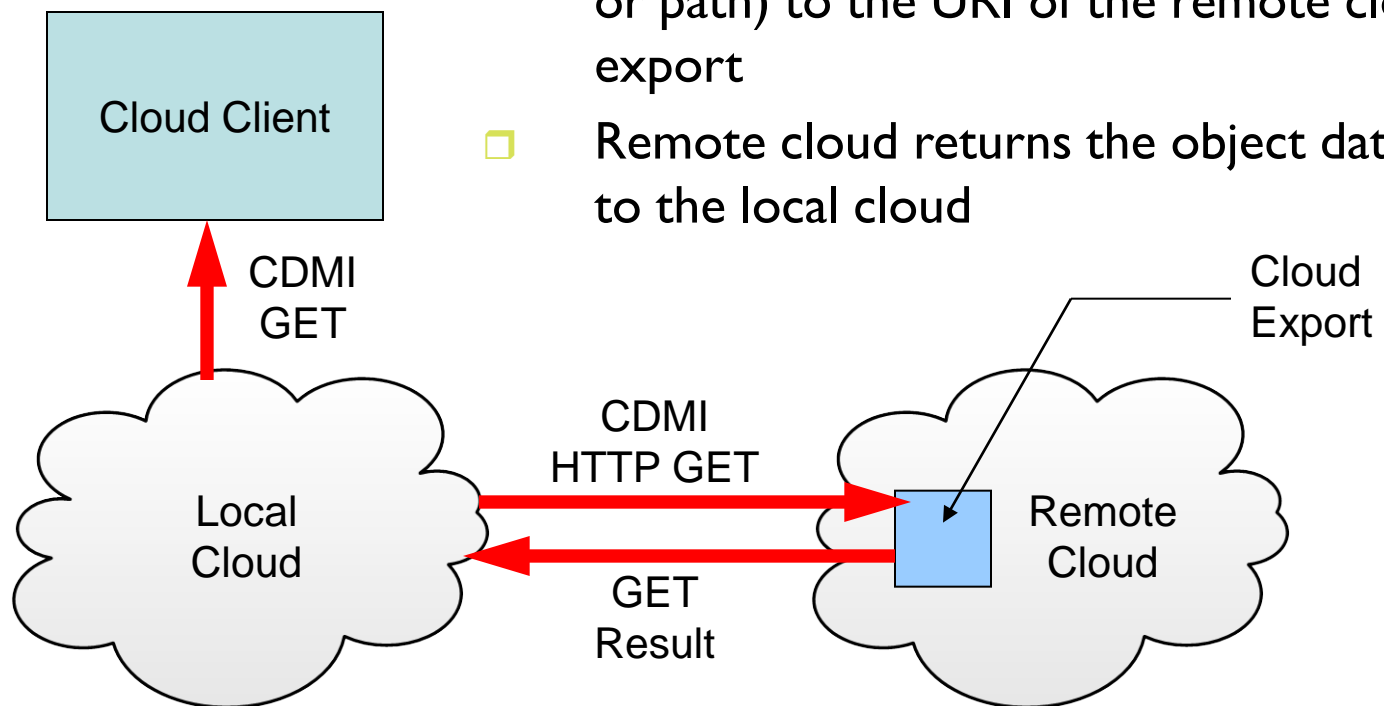




# Storage Federation (Level 2)


2. If a client requests the object, but it is only stored remotely:

- Local cloud does an CDMI GET (by ID or path) to the URI of the remote cloud export
- Remote cloud returns the object data to the local cloud




# Storage Federation (Level 2)

- Local cloud can create a CDMI notification queue in the domain URI on the remote cloud to receive notifications about events related to objects in the domain:
  - Local cloud does an CDMI GET (by ID or path) to the URI of the remote notification queue
  - Remote cloud returns the events related to stored objects



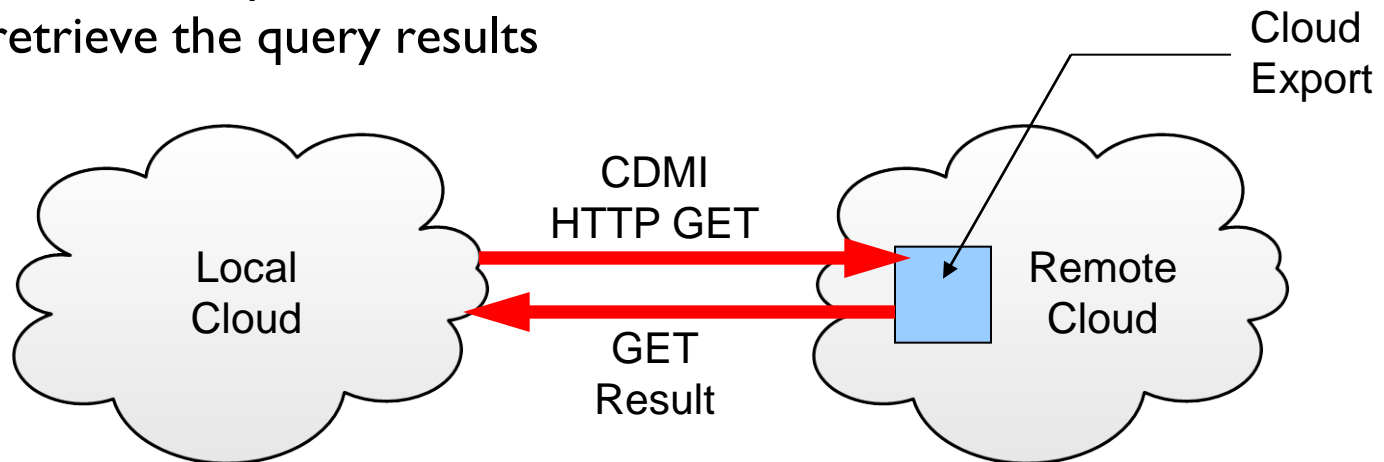
# Storage Federation (Level 2)

- Local cloud can create a CDMI logging queue in the domain URI on the remote cloud to receive logs about events related to objects in the domain:
  - Local cloud does an CDMI GET (by ID or path) to the URI of the remote logging queue
  - Remote cloud returns the logs related to stored objects




# Storage Federation (Level 2)

5. Local cloud can create a CDMI query queue in the domain URI on the remote cloud to perform queries across objects stored on the remote cloud:
  - ❑ Local cloud performs a CDMI PUT (by ID or path) to the URI of the remote query queue, indicating where the query results should be returned
  - ❑ Local cloud performs a CDMI GET to retrieve the query results



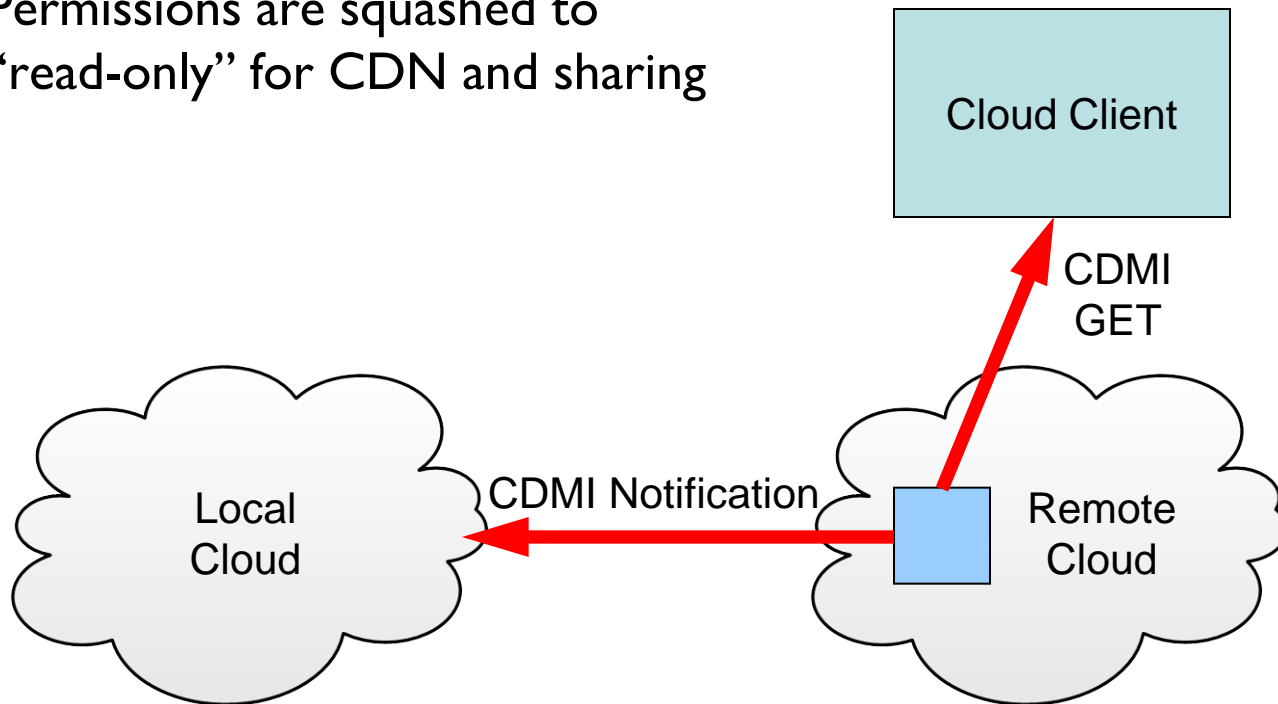
# Access Federation (Level 3)

- I. In level 2 federation, only the local cloud has permissions to access the cloud export. In level 3 federation, the domain membership on the remote cloud contains client entries as well as the credentials for the local cloud.
  - Domain membership may be delegated to a common source (eg, AD) or back to the local cloud



# Access Federation (Level 3)

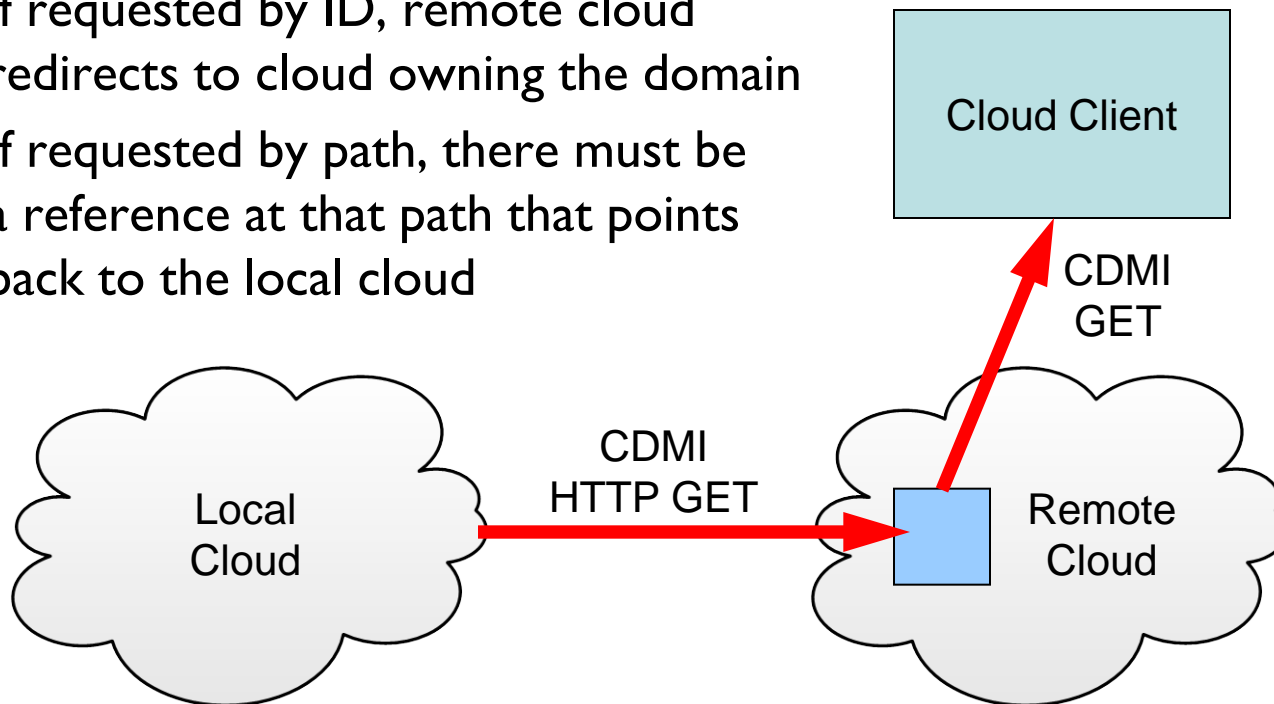
2. When a client accesses an object directly from the remote cloud, their credentials are looked up in the domain membership.
  - Permissions are squashed to “read-only” for CDN and sharing



# Access Federation (Level 3)


3. If a client accesses an object that is only stored on the local cloud (either by CDMI ID or by path), the remote cloud retrieves the object from the local cloud.

- If requested by ID, remote cloud redirects to cloud owning the domain
- If requested by path, there must be a reference at that path that points back to the local cloud



# Peering Federation (Level 4)

- I. In level 3 federation, additions and changes to the stored content can only be done through the local cloud. In level 4 federation, changes can be performed on either clouds.
  - ❑ Domain entry delegates back to the local cloud
  - ❑ Local cloud must monitor remote cloud using notifications

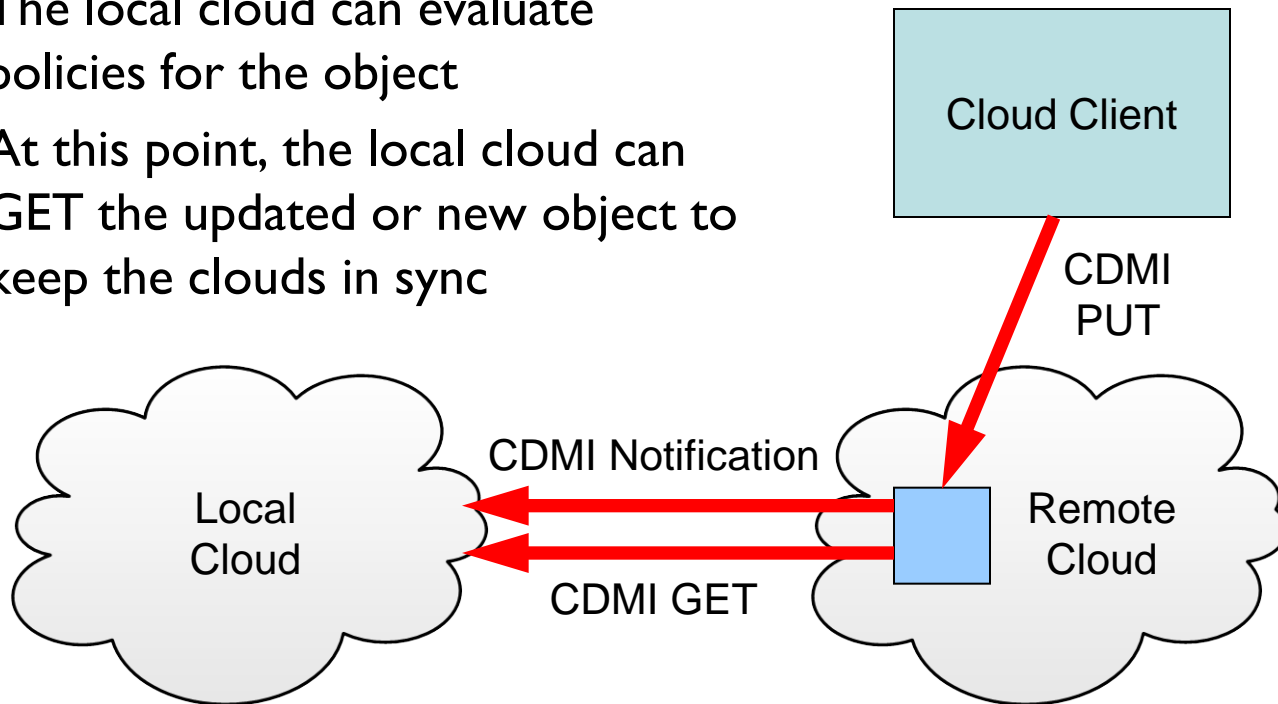




# Peering Federation (Level 4)

2. When a client modifies or creates an object on the remote cloud in the CDMI domain being managed by the local cloud, the local cloud receives a notification.

- The local cloud can evaluate policies for the object
- At this point, the local cloud can GET the updated or new object to keep the clouds in sync



- ❑ CDMI Identifiers enable Federation
  - ❑ Because OIDs are location-independent, once federated, a request for an object that is not stored locally can be redirected to the originating domain.
  - ❑ Path objects can be likewise redirected if path references are created

- ❑ CDMI Domains Enable Federation
  - ❑ Every object belongs to one and only one domain
  - ❑ Domains indicate which content is being managed by which cloud
  - ❑ There is only one cloud that manages content belonging to a domain
  - ❑ In order to federate two clouds, there must be an identically-named domain on both clouds

- ❑ While CDMI features enable federation, CDMI does not explicitly specify how federation works
  - ❑ Possible new work item for the CDMI Technical Working Group to identify what features are required for federation
  - ❑ Possible new work item to define a set of interoperability tests to facilitate inter-vendor federation

# Thank you!

## Questions and Answers

Contact Info:  
[dslik@netapp.com](mailto:dslik@netapp.com)